Yet another alternative approach to a scoring function would be one that computes the probability that a given potential defect represents an actual defect. This is perhaps the most elegant method possible. This might be reasonable in some applications where the images acquired have predictable properties. Such applications might include astronomy, where the density of stars may be largely predictable. If the images are of manufactured object in which defects are closely controlled, this might also be possible. In this case, the score target value h would just be the demanded probability, for instance 1.0 minus 1 millionth, that a detected defect be an actual defect.

8.12 Defect Removal Policy

In order to maintain the map of potential defects to a reasonable size, a policy for removing some of the regions from it must be applied in step **840** of FIG. **8**. The simplest policy is to remove regions that arc a certain number of qualified images out of date—that is, came from images acquired more than a predetermined number z images ago.

However, in some cases there may be superior polices. More generally, if we are demanding of a potential defect r a count of k hits in the most recent n images in which r is active, then a region is obsolete once it is active in n images more recent than the image from which the region was derived.

Instituting this policy provides a basic limit on the size of potential defects. Suppose our images have a number of pixels denoted by A and that we use the optional step of making inactive segments that have a total deviation above a certain threshold t. Then in the worst case $t \cdot A$ potential defects must be stored. If we assume the space taken up by each defect in the map is at worst a constant c for a single pixel region, then maximum space required for defects is $c \cdot t \cdot A \cdot n$. Reasonable values for these variables might be an image area of 100 million pixels, a threshold t of 1/1000, a value for n of 5 images, and a value for c of 20 bytes. The potential defect map would at most take up 10 million bytes in this case. At the time of this writing that seems extravagant, but can be expected to become trivial in a few years. A wide variety of applications exists today which are much easier: CCD images have closer to 4 million pixels, and a threshold t of 1/10,000 would work fine for clean white printing paper. Also, deviant regions being individual pixel is the worst case for this analysis, as regions shapes can be stored much more compactly than a similar sized group of individual pixels, so c might be much less than 1.0 in practice.

This allows us to set aside, either in a general purpose computer or in a memory that is part of an image acquistion apparatus, a fixed storage that will not be exceeded by this algorithm. This is another indicaion that the "k of n" embodiment is best.

9 Conclusion, Ramifications and Scope

This defect detection method allows automatic dynamic detection of dirt, dust, and scratches on the glass of photocopiers and scanners or defects in the CCD arrays of video cameras. No special action is required of the operator, and no special object, such as an unblemished target, is required. It thus simplifies the operation and improves the images produced by such machines by reliably detecting and reporting such defects.

While the above description contains many specifics for the sake of clarity, it should not be construed as limitations on the scope of the defect detection method, but rather as an exemplification of several preferred embodiments. In particular, once a scoring function is chosen equivalent optimized embodiments become straightforward. More generally, a wide applicability to instruments and situations is possible.

For example, the defect identification method can be usefully employed in video cameras. In such an instance, a majority of the images captured would not yield any sufficiently homogenous segments to be counted as active. However, when indoors, simply pointing the camera occasionally at a typical relatively smooth wall would suffice to provide the method a sufficient number of generally homogeneous images to identify defects. This process could occur passively, or the operator could choose to identify defects by intentionally aiming the camera at such images. Out-of-doors, the sky provides a convenient generally homogeneous background, even if starry or partly cloudy.

CCD arrays such as those used in video cameras are also employed in telescopes and other scientific instruments, and hence the defect identification method can be employed on such devices as well. Since defects are disastrous in such instruments and the scenes or samples they target are very easily controlled, the defect identification method may be very useful in such instruments. Similar arrays of defect-susceptible detectors are also used for infrared imaging.

Accordingly, the scope of the invention should be determined not by the embodiments illustrated, but by the appended claims and their legal equivalents.

9.1 Limitations on the Usefulness of the Defect Detection Procedure

The strengths of a technology can often best be understood by considering its weaknesses. Accordingly, this section discusses what the defect detection procedure cannot do.

Because it relies on probability, it can only assure an extremely low probability of falsely detecting a defect, rather than eliminating all such possibility. Because it requires the comparison of many images, it cannot detect a defect instantly, but only after several qualified images have been acquired. It follows that a defect that comes and goes quickly, such as a dust most that is laid down in one image and wiped away by the next, cannot be detected. Similarly, a defect cannot be detected unless it falls in a generally homogeneous area in a series of images. Thus, if the images are always "busy" at a certain spot, that is, always containing text or edges, then defects in that region cannot be detected. For instance, all U.S. Patents contain nearly the same text in the top margin of each page. The photocopying of a U.S. Patent would not be able to detect defects in that region, but would be delayed until by chance documents that are blank in that area are scanned. Similarly, the dense text of the a patent would not allow defect detection in that area on those pages, though the Sheets on which drawings appear might well contain enough blank spaces to allow defect detection in those areas.

In certain simple embodiments the method can detect large defects. However, if the option of segmenting the image into inactive and active samples is used either by the total deviation method or by the text segmentation method, then only relatively small defects can be detected. Large defects will in such circumstances generally be considered inactive segments and therefore no defects will be found therein. Of course, large defects are catastrophic and clearly noticeable, so detecting them is not as important as detecting the small speckles which plague photocopiers and scanners.